arXiv:1612.04694v2 [math.OC] 19 Dec 2016

Stochastic Second-Order Optimization via von Neumann Series

Mojmír Mutný
Department of Mathematics
ETH Zürich
Zürich, Switzerland
mmutny@student.ethz.ch

Abstract

A stochastic iterative algorithm approximating second-order information using von Neumann series is discussed. We present an improved and simplified convergence analysis, in contrast to a similar algorithm and its analysis, LISSA. The algorithm is primarily suitable for training large scale linear models, where the number of data points is very large. Two novel analyses, one showing space independent linear convergence, and one showing conditional quadratic convergence are discussed. In numerical experiments, the behavior of the error is similar to the second-order algorithm L-BFGS, and it performs better than a similar algorithm LISSA.

1 Introduction

With the advent of the age of Big Data, the need for novel, fast and robust optimization algorithms becomes ever more prevalent. The two most common classes of optimization algorithms are first-order methods which utilize only gradient information, and on the other hand second-order methods, which incorporate curvature information from the objective function. In this paper, we explore the possibility of using the von Neumann approximation series as a tool to iteratively build up second-order information for efficient optimization methods.

Many supervised learning problems can be formulated as empirical risk minimization,

$$F(x) := \frac{1}{n}\sum_{i=1}^{n} f(x, z^i, y^i). \tag{1}$$

where $\{z^i, y^i\}$ are data points lying in $\mathbb{R}^d$, with associated labels $y^i \in \mathbb{R}$. The optimization variable $x \in \mathbb{R}^d$ is the parametrization of the model that one seeks to determine by solving the following unconstrained minimization problem,

$$x^* := \arg\min_{x\in\mathbb{R}^d} F(x). \tag{2}$$

We focus on the scenario when the number of data points is much larger than the dimension of the problem i.e. $n \gg d$. In this scenario, the methods that belong to standard optimization routines are SGD [2], dual SDCA [14] or more advanced algorithms such as SVRG [5], SAGA A[4], SAG [13], or S2GD [6], which collectively belong to the school of first-order methods. In Big Data optimization, the school of second-order optimization is represented by L-BFGS [7] algorithm, and many of its varieties. Apart from L-BFGS, ideas of sub-sampling the Hessian recently formed a plethora of algorithm in literature [12, 11, 10, 8].

We propose a method, where access to unbiased estimates of a Hessian $\{\mathbf{X}_i\}_{i=1}^n$ is assumed. These estimates can be used to iteratively build second-order information exploiting the well known von

Neumann series approximation (3). In fact, our method builds asymptotically inverse Hessian in the expectation. One can contrast this sampling with SGD, where gradients are sampled, and in expectation true gradient is built. In our analysis, we assume exact gradient information and the Hessian matrix is sampled. However, an extension to inexact first-order information could be feasible as well.

Remark 1 ([3])**.** *For a matrix $\mathbf{A} \in \mathbb{R}^{d\times d}$ such that $\mathbf{A} \succeq 0$ and $\|\mathbf{A}\|_2 \leq 1$ we have that*

$$\mathbf{A}^{-1} = \sum_{j=0}^{\infty} (\mathbf{I} - \mathbf{A})^j. \tag{3}$$

1.1 LISSA algorithm

A very similar idea to ours has been presented in [1] with slight differences in the algorithm formulation. The main focus of [1] was on averaging the estimators of Hessian to gain better estimation properties utilizing the theory of concentration inequalities. Their analysis of the algorithm requires warm start by gradient descent to prove linear convergence to the optimum. The algorithm LISSA from [1] will not be reviewed here, but essentially it consist of two parts parametrized by two parameters, S_1 and S_2. The parameter S_2 denotes the number of Hessian samples used to build one estimate and S_1 denotes number of repetitions of this step. In numerical experiments, however, authors set $S_1 = 1$ corresponding to absence of averaging.

1.2 Contributions

In this work, we present two new simplified convergence analyses of the stochastic second-order algorithm that builds approximation of a inverse Hessian using unbiased samples of a Hessian matrix. We show more intuitive convergence rates than a similar algorithm [1] under milder mathematical conditions. Namely, we do not require initial convergence to a closed proximity of an optimum to show a linear convergence, and we should conditional second order algorithm given a specific condition on gradient. Our formulation of algorithm does not require notion of averaging, and thus reduces number of parameters needed to perform the theoretical analysis. Additionally, the theoretical formulation of our algorithm is more closely aligned with its practical execution.

2 Algorithm

In order to define the algorithm formally, we first have to introduce a concept of an index sampling i.e. set-valued random variable. We closely follow the notation and theory of sampling of [9]. Furthermore, we propose our first two crucial assumptions. The Assumption 1 might seem unnatural but we note that most of the problems can be recast such that it holds.

Definition 1 (ordered τ-independent sampling)**.** *An ordered τ-independent sampling $\hat{S}$ is ordered random set-valued mapping with elements from $[n]$ and cardinality τ. In addition, the possible outcomes of this random variance are all equally probable.*

Assumption 1 (Hessian)**.** *The function $f(x) : \mathbb{R}^d \to \mathbb{R}$ is twice continuously differentiable, and has scaled Hessian s.t. $\nabla^2 f(x) \preceq \mathbf{I}\ \forall x \in \mathbb{R}^d$.*

Assumption 2 (Hessian Samples)**.** *We assume existence of $\{\mathbf{X}_i\}_{i=1}^n$, with each $\mathbf{X}_i \in \mathbb{R}^{d\times d}$, at each point x where each $\mathbb{E}[\mathbf{X}_i] = \nabla^2 f(x)$ for a given x. To signify the x dependence, we denote these samples as*

$$\tilde{\nabla}^2 f^j(x) := \mathbf{X}_j. \tag{4}$$

The Algorithm is presented in 1. The step-size is fixed before. The usual parameter used is $c = 1$. Additionally, we present two versions of the algorithm, *practical* and *theoretical*. If Hessians samples are independent of x^k the *practical* version of the algorithm can be used, where previous estimates of Hessian can be reused. In the other more general case, a full estimator has to be rebuilt from the start, which can be grossly inefficient. In such cases, the series must be truncated for practical purposes.

A concrete example of the empirical risk minimization problem which fulfills Assumptions 1 and 2 is ridge regression fitting with a cost function $\frac{1}{2n}\sum_{i=1}^n (z_i^\top x - y_i)^2 + \frac{\lambda}{2}\|x\|_2$. The Hessian can

Algorithm 1 Iterative Stochastic Second-Order Algorithm (ISSA)

Parameters: Sampling $\hat{S}$, $c \in \mathbb{R}$ according to the Theorems 1 and 2
Initialization: Pick $x_0 \in \mathbb{R}^d$, $\mathbf{A}_0 = \mathbf{I}$, $S^0 = \{\}$

for $k = 0, 1, 2, \ldots$ **do**
  Sample $\hat{S}$ to get $S^k = S^{k-1} \cup \hat{S}$
  if Hessian independent of x^k **then**
    $\mathbf{A}_k = \mathbf{A}_{k-1}$
    for $j \in \hat{S}$ **do**
      $\mathbf{A}_k = \mathbf{I} + (\mathbf{I} - \tilde{\nabla}^2 f^j)\mathbf{A}_k$
    end for
  else
    $\mathbf{A}_k = \mathbf{I}$
    for $j \in S^k$ **do**
      $\mathbf{A}_k = \mathbf{I} + (\mathbf{I} - \tilde{\nabla}^2 f^j(x^k))\mathbf{A}_k$
    end for
  end if
  $\Delta x = \mathbf{A}_k \nabla f(x^k)$
  $x^{t+1} \leftarrow x^t - \frac{1}{c}\Delta x$
end for

be easily calculated $\nabla^2 f(x) = \left(\frac{1}{n}\sum_{i=1}^n z_i z_i^\top + \lambda \mathbf{I}\right)$, and the Hessian samples enumerated with $i \in [n]$ can be defined as $\mathbf{X}_i :\overset{\texttt{def}}{=} z_i z_i^\top + \mathbf{I}\lambda$.

3 Convergence Analysis

3.1 Assumptions

Assumption 3 (Smoothness, Strong Convexity)**.** *There exists a positive constant α, β such that $\forall x, h \in \mathbb{R}^d$,*

$$f(x) + \langle \nabla f(x), h\rangle + \frac{\alpha}{2}\|h\|_2^2 \le f(x+h) \le f(x) + \langle \nabla f(x), h\rangle + \frac{\beta}{2}\|h\|_2^2 . \tag{5}$$

Minimizing the above on both sides in h gives:

$$f(x) - f(x^*) \le \frac{1}{2\alpha}\langle \nabla f(x), \nabla f(x)\rangle \tag{6}$$

where x^* denotes the minimum point of f.

3.2 Matrix Notation

One of the contributions of our work is introduction of a new notation to prove the convergence results presented in Definition 2.

Definition 2 (Matrix Definition of the update)**.** *Let $\mathbf{X}_i$ be a matrix valued random variable such that $\mathbb{E}[\mathbf{X_i}] = \nabla f^2(x)$ and $\hat{S}$ be an ordered τ-independent sampling such that $|S| = k\tau$, then define*

$$\mathbf{Y}_i := \begin{pmatrix} (\mathbf{I} - \mathbf{X}_i) & \mathbf{I} \\ \mathbf{0} & \mathbf{I} \end{pmatrix} \qquad \textit{and,} \qquad \mathbf{R}^k := (\mathbf{I}\,\mathbf{0}) \left(\prod_{j \in S} \mathbf{Y}_j\right) (\mathbf{I}\,\mathbf{I})^\top .$$

Lemma 1. *Let us have the same assumptions as in Definition 2,*

$$\left\|\mathbb{E}[\mathbf{R}^k] - (\nabla^2 f(x))^{-1}\right\|_2^2 \le \frac{(1-\alpha)^{k\tau}}{\alpha} \tag{7}$$

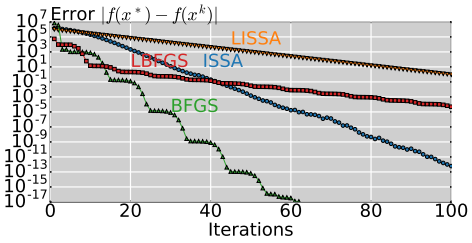

(a) A linear regression model with an artificial dataset such that $n = 10^6$ and $m = 10^2$.

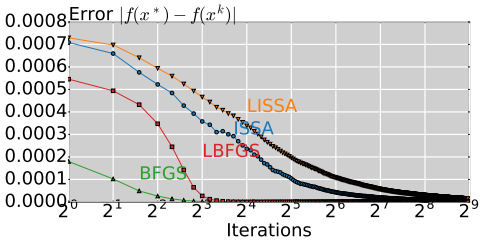

(b) A logistic regression on *mushroom* dataset with $n = 8124$ and $m = 22$.

Figure 1: A numerical experiment comparing ISSA, LISSA, BFGS and L-BFGS algorithms. In a) the $\tau = 5$ for ISSA and L-BFGS used last 5 gradient information as well. LISSA was run with $S_2 = 20, S_1 = 1$. In the second experiment on the real dataset all parameters were equal to 3. ISSA was run without the decreasing step size.

3.3 Linear Convergence

Theorem 1. *Let Assumptions 1, 2 and 3 be satisfied then the algorithm converges to the optimum denoted x^* with $c = \beta\left(\frac{(2-\alpha^2)}{\alpha} + (1-\alpha)^k\right)$ linearly in expectation with rate,*

$$\mu := \left(\frac{\alpha^2}{\beta(2-\alpha^2)}\right) \tag{8}$$

$$\mathbb{E}[f(x^{k+1}) - f(x^k)] \leq -\mu(f(x^k) - f(x^*)) \tag{9}$$

Note than in contrast to the results of [1], we prove linear convergence without the need to converge to sufficient vicinity of the optimum. However, the result is much worse than a standard gradient descend with rate $\frac{\alpha}{\beta}$, and is only the consequence of the Hessian scaling, or Hessian correction not interfering to such extent as to hinder the optimization.

3.4 Quadratic Convergence

Theorem 2. *Let Assumptions 1, 2 and 3 be satisfied, and additionally let*

$$\frac{\alpha^2}{8\beta} \geq \left\|\nabla f(x^k)\right\| \geq \frac{\beta(1-\alpha)^{k\tau}}{2} + \frac{\alpha\beta(\beta-\alpha)}{4}, \tag{10}$$

then the algorithm improves the number of significant digits quadratically when run with $c = 1$.

We would like to remark that the last condition for Theorem 2 is rather artificial, and is hardly satisfied for the whole duration of the optimization; if at all. However, empirical results from [1], as well as ours, show a very steep initial convergence to the optimum, which might hint that in such circumstances the conditions of the theorem is satisfied and is the consequence of the initial rapid decrease. For example, consider, τ large , $\beta = 1$ and $\alpha \geq 1/3$, then we see that a suitably chosen starting point could lead to gradient fulfilling the condition allowing for very fast convergence.

4 Numerical Results

The experiment was performed on Lenovo Laptop with processor Intel i7 2.6 GHz. We used the numerical linear algebra library Eigen as the main backbone for our implementation. We show one experiment fitting least squares estimator in Figure 1 with artificial dataset that where the design matrix elements are sampled from truncated standard normal distribution, and logistic regression on *mushroom* dataset 1.

5 Conclusion

We presented an iterative stochastic second-order optimization algorithm that iteratively builds an approximation to Hessian inverse via Neumann Series. We have provided two analyses of the conver-

gence properties of the algorithm. One of those shows linear convergence under milder conditions than previous results and one analysis which might shed light on the very fast initial convergence of the method. The numerical experiment shows that our method matches L-BFGS performance.

Acknowledgments

The author would like to thank Prof Nicolai Meinshausen of ETH Zürich, and Prof Martin Jaggi of EPF Lausanne for their kind advice and help with this project.

References

[1] Naman Agarwal, Brian Bullins, and Elad Hazan. Second order stochastic optimization in linear time. *arXiv preprint arXiv:1602.03943*, 2016.

[2] Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pages 177–186. Springer, 2010.

[3] Ke Chen. *Matrix preconditioning techniques and applications*, volume 19. Cambridge University Press, 2005.

[4] Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. In *Advances in Neural Information Processing Systems*, pages 1646–1654, 2014.

[5] Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. In *Advances in Neural Information Processing Systems*, pages 315–323, 2013.

[6] Jakub Konecnỳ, Zheng Qu, and Peter Richtárik. S2cd: Semi-stochastic coordinate descent. In *NIPS Optimization in Machine Learning workshop*, 2014.

[7] Dong C Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical programming*, 45(1-3):503–528, 1989.

[8] Mert Pilanci and Martin J Wainwright. Newton sketch: A linear-time optimization algorithm with linear-quadratic convergence. *arXiv preprint arXiv:1505.02250*, 2015.

[9] Zheng Qu and Peter Richtárik. Coordinate descent with arbitrary sampling ii: Expected separable overapproximation. *arXiv preprint arXiv:1412.8063*, 2014.

[10] Zheng Qu, Peter Richtárik, Martin Takáč, and Olivier Fercoq. Sdna: stochastic dual newton ascent for empirical risk minimization. *arXiv preprint arXiv:1502.02268*, 2015.

[11] Farbod Roosta-Khorasani and Michael W Mahoney. Sub-sampled newton methods i: Globally convergent algorithms. *arXiv preprint arXiv:1601.04737*, 2016.

[12] Farbod Roosta-Khorasani and Michael W Mahoney. Sub-sampled newton methods ii: Local convergence rates. *arXiv preprint arXiv:1601.04738*, 2016.

[13] Mark Schmidt, Nicolas Le Roux, and Francis Bach. Minimizing finite sums with the stochastic average gradient. *arXiv preprint arXiv:1309.2388*, 2013.

[14] Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. *Journal of Machine Learning Research*, 14(Feb):567–599, 2013.

Appendix

Lemma 2. *Let $\mathbf{R}_k$ be defined as in Definition 2. Then the following bound on the second moment holds,*

$$\lambda_{max}(\mathbb{E}[(\mathbf{R}_k)^\top \mathbf{R}_k]) \leq \frac{2-\alpha^2}{\alpha} + (1-\alpha)^k. \tag{11}$$

Proof. We use shorthand for $v_{k+1} = \lambda_{max}(\mathbb{E}[(\mathbf{R}_k)^\top \mathbf{R}_k])$, $c = 1 + \frac{2}{\alpha}$ and $\sigma_{k+1} = v_{k+1} + \frac{2}{\alpha}$. Also, we know that $v_1 = 1$. Then using the definition of update in Algorithm 1, we get,

$$
\begin{aligned}
(\mathbf{R}_k)^\top \mathbf{R}_k &= (\mathbf{I} + \mathbf{R}_k^\top (\mathbf{I} - \mathbf{X}_j)^\top)(\mathbf{I} + (\mathbf{I} - \mathbf{X}_j)\mathbf{R}_k) && (12)\\
&= \mathbf{I} + \mathbf{R}_k^\top (\mathbf{I} - \mathbf{X}_j)^\top + (\mathbf{I} - \mathbf{X}_j)\mathbf{R}_k + \mathbf{R}_k^\top (\mathbf{I} - \mathbf{X}_j)^2 \mathbf{R}_k && (13)\\
&\overset{Ass1}{\leq} \mathbf{I} + \mathbf{R}_k^\top (\mathbf{I} - \mathbf{X}_j)^\top + (\mathbf{I} - \mathbf{X}_j)\mathbf{R}_k + \mathbf{R}_k^\top (\mathbf{I} - \mathbf{X}_j)\mathbf{R}_k && (14)\\
v_{k+1} &\leq 1 + (1-\alpha)\left(\frac{2}{\alpha} + v_k\right) && (15)\\
v_{k+1} + \frac{2}{\alpha} &\leq 1 + \frac{2}{\alpha} + (1-\alpha)\left(\frac{2}{\alpha} + v_k\right) && (16)\\
\sigma_{k+1} &\leq c + (1-\alpha)\sigma_k && (17)\\
&\leq c \sum_{i=0}^{k-1} (1-\alpha)^i + (1-\alpha)^k && (18)\\
&\leq \frac{c}{\alpha} + (1-\alpha)^k && (19)
\end{aligned}
$$

We remark that given k large it is reasonable to assume stationary and we can check that the calculated bound, assuming stationarity from (14) matches the non-stationary result. □

Proof of Theorem 1. We adopt the notation shorthand that $\nabla f(x^k) = g^k$, and $c = \beta\left(\frac{(2-\alpha^2)}{\alpha} + (1-\alpha)^k\right)$,

$$
\begin{aligned}
f(x^{k+1}) - f(x^k) &\overset{(5)}{\leq} -\frac{1}{c}\langle g^k, \mathbf{R}^k g^k\rangle + \frac{\beta}{2c^2}\langle \mathbf{R}^k g^k, \mathbf{R}^k g^k\rangle\\
\mathbb{E}[f(x^{k+1}) - f(x^k)] &\leq -\frac{1}{c}\langle g^k, \mathbb{E}[\mathbf{R}^k] g^k\rangle + \frac{\beta}{2c^2}\langle g^k, \mathbb{E}[(\mathbf{R}^k)^2] g^k\rangle\\
&\overset{(11)}{\leq} -\frac{1}{c}\langle g^k, g^k\rangle + \frac{\beta}{2c^2}\left(\frac{(2-\alpha^2)}{\alpha} + (1-\alpha)^k\right)\langle g^k, g^k\rangle\\
&\overset{(c)}{\leq} -\frac{\alpha}{2\beta(2-\alpha^2)}\langle g^k, g^k\rangle\\
&\overset{(5)}{\leq} -\frac{\alpha^2}{\beta(2-\alpha^2)}(f(x^k) - f(x^*))\\
&\overset{(8)}{\leq} -\mu(f(x^k) - f(x^*))
\end{aligned}
$$

□

Lemma 3. *Let us have the same assumptions as in Definition 2,*

$$\mathbb{E}[\mathbf{R}^k] = \sum_{j=1}^{k\tau} (\mathbf{I} - \nabla^2 f(x^k))^j. \tag{20}$$

Proof of Lemma 1 . Using von Neumann expansion

$$
\begin{aligned}
\left\|\mathbb{E}[\mathbf{R}^k] - (\nabla^2 f(x))^{-1}\right\|_2^2 &\overset{(20)}{=} \left\|\sum_{j=k}^{\infty} (\mathbf{I} - \nabla^2 f(x))^{-1}\right\| && (21)\\
&\leq (1-\alpha)^{\tau k} \sum_{k=0}^{\infty} \left\|(\mathbf{I} - \nabla^2 f(x))^{-1}\right\|_2 && (22)\\
&= \frac{(1-\alpha)^{k\tau}}{\alpha} && (23)
\end{aligned}
$$

□

Proof of Theorem 2.

$$\Delta x^k := -\mathbf{R}^k \nabla f(x^k) \tag{24}$$

$$\begin{aligned}
\left\|\nabla f(x^{k+1})\right\| &= \left\|\nabla f(x^{k+1}) + \nabla f(x^k) - (\mathbf{R}^k)^{-1}\Delta x^k\right\| \\
&\leq \left\|\nabla f(x^{k+1}) + \nabla f(x^k) + \left(\nabla^2 f(x^k) - \nabla^2 f(x^k) - (\mathbf{R}^k)^{-1}\right)\Delta x^k\right\| \\
&\leq \left\|\nabla f(x^{k+1}) + \nabla f(x^k) - \nabla^2 f(x^k)\Delta x^k\right\| + \left\|\left(\nabla^2 f(x^k) - (\mathbf{R}^k)^{-1}\right)\Delta x^k\right\| \\
&\leq 2\beta\left\|\Delta x^k\right\|^2 + \left\|(\mathbf{R}^k)^{-1}\right\|\left\|\nabla^2 f(x^k)\right\|\left\|\nabla^2 f(x^k)^{-1} - \mathbf{R}^k\right\|\left\|\Delta x^k\right\| \\
&\leq 2\beta\left\|\Delta x^k\right\|^2 + \beta^2\left\|\nabla^2 f(x^k)^{-1} - \mathbf{R}^k\right\|\left\|\Delta x^k\right\| \\
&\leq 2\beta\left\|\Delta x^k\right\|^2 + \beta^2\left(\left\|\nabla^2 f(x^k)^{-1} - \mathbb{E}[\mathbf{R}^k]\right\| + \left\|\mathbb{E}[\mathbf{R}^k] - \mathbf{R}^k\right\|\right)\left\|\Delta x^k\right\| \\
&\leq 2\frac{\beta}{\alpha^2}\left\|\nabla f(x^k)\right\|^2 + \frac{\beta^2}{\alpha}\left(\left\|\nabla^2 f(x^k)^{-1} - \mathbb{E}[\mathbf{R}^k]\right\| + \left\|\mathbb{E}[\mathbf{R}^k] - \mathbf{R}^k\right\|\right)\left\|\nabla f(x^k)\right\| \\
\mathbb{E}[\left\|\nabla f(x^{k+1})\right\|] &\overset{(7)}{\leq} 2\frac{\beta}{\alpha^2}\left\|\nabla f(x^k)\right\|^2 + \frac{\beta^2}{\alpha}\left(\frac{(1-\alpha)^{k\tau}}{\alpha} + \frac{\beta-\alpha}{2}\right)\left\|\nabla f(x^k)\right\| \\
&= 2\frac{\beta}{\alpha^2}\left\|\nabla f(x^k)\right\|\left(\left\|\nabla f(x^k)\right\| + \frac{\beta(1-\alpha)^{k\tau}}{2} + \frac{\alpha\beta(\beta-\alpha)}{4}\right) \\
&\overset{(10)}{\leq} 4\frac{\beta}{\alpha^2}\left\|\nabla f(x^k)\right\|^2
\end{aligned}$$

Lastly, (6) implies,

$$\mathbb{E}[f(x^k) - f(x^l)] \leq \frac{\alpha}{8\beta}\left(\frac{1}{2}\right)^{2^{(k-l)}}, \tag{25}$$

where l is the index such that the (10) still holds. □